arXiv:1403.7429v1 [math.OC] 28 Mar 2014

Distributed Reconstruction of Nonlinear Networks: An ADMM Approach ⋆

Wei Pan, Aivar Sootla, and Guy-Bart Stan

Centre for Synthetic Biology and Innovation and the Department of Bioengineering, Imperial College London, United Kingdom; e-mail: {w.pan11, a.sootla, g.stan}@ imperial.ac.uk

Abstract: In this paper, we present a distributed algorithm for the reconstruction of large-scale nonlinear networks. In particular, we focus on the identification from time-series data of the nonlinear functional forms and associated parameters of large-scale nonlinear networks. In (Pan et al. (2013)), a nonlinear network reconstruction problem was formulated as a nonconvex optimisation problem based on the combination of a marginal likelihood maximisation procedure with sparsity inducing priors. Using a convex-concave procedure (CCCP), an iterative reweighted lasso algorithm was derived to solve the initial nonconvex optimisation problem. By exploiting the structure of the objective function of this reweighted lasso algorithm, a distributed algorithm can be designed. To this end, we apply the alternating direction method of multipliers (ADMM) to decompose the original problem into several subproblems. To illustrate the effectiveness of the proposed methods, we use our approach to identify a network of interconnected Kuramoto oscillators with different network sizes (500∼100,000 nodes).

1. INTRODUCTION

The importance of reconstructing nonlinear systems and its associated difficulties are widely recognised (Ljung et al. (2011)). Reconstruction methods focus on specific system classes such as those described by Wiener and Volterra series or nonlinear auto-regressive with exogenous inputs (NARX) models to name just a few examples (see Ljung (1999) and references therein). However, nonlinear systems can be described by other functional forms. One of the most important and challenging problems in nonlinear network reconstruction is nonlinear structure identification (Sjöberg et al. (1995)). Nonlinear functional forms can be typically expanded as sums of terms belonging to a family of parameterised functions (see Sec. 5.4, Ljung (1999)). A usual approach to identifying a nonlinear black-box model is to search amongst a set of possible nonlinear terms (e.g., basis functions) for a parsimonious description coherent with the available data set Haber and Unbehauen (1990)).

In this paper, nonlinear systems are represented in a general state-space form. The framework we develop uses some a priori knowledge of the type of system we want to reconstruct, i.e., we consider a set of candidate dictionary functions appropriate for the specific type of systems from which the data have been collected (e.g., biological, chemical, mechanical, or electrical system). We assume the measured time-series data and a set of candidate dictionary functions are given. Our main objective is to identify the most parsimonious representation that explains the collected time-series data at best. Generally, we cast the reconstruction problem into a sparse signal recovery problem (Candès and Tao (2005); Donoho (2006)). In Pan et al. (2013), the nonlinear network reconstruction problem was casted as a nonconvex optimisation problem which was shown to be efficiently solvable using a centralised reweighted lasso algorithm. Nonlinear reconstruction problems solved by centralised reconstruction methods have typically a relatively small size.

Social networks, communication networks and biological networks are typically very large (e.g. more than 100,000 nodes) and the data set collected from them is therefore quite "big". Typically, centralised reconstruction algorithms cannot handle such problems due to their associated very large memory and computational requirements. Here, we will apply the alternating direction method of multipliers (ADMM) to split the centralised problem into several subproblems with each subproblem solving a weighed lasso problem independently. This approach has the advantage that memory and computational requirements can be both reduced in comparison to generic centralised solvers.

ADMM is a powerful algorithm for solving structured convex optimization problems. The ADMM method was introduced for optimisation in the 1970's and is closely related to many other optimisation algorithms including Bregman iterative algorithms, Douglas-Rachford splitting, and proximal point methods (see (Boyd et al. (2011)) and references therein). ADMM has been shown to have strong convergence properties and to be useful for solving by decomposition large optimisation problems which cannot be handled by generic optimization solvers. ADMM has been applied in many areas, such as filtering (Wahlberg et al. (2012)), image processing (Figueiredo and Bioucas-Dias (2010)) as well as large-scale problems in statistics and machine learning (Boyd et al. (2011)).

⋆ W. Pan gratefully acknowledge the support of Microsoft Research through the PhD Scholarship Program. A. Sootla and G.-B. Stan acknowledge the support of EPSRC through the project EP/J014214/1 and the EPSRC Science and Innovation Award EP/G036004/1.

The paper is organised as follows. In Section 2, we formulate the nonlinear network reconstruction problem we consider in this paper. In Section 3, we re-interpret this reconstruction problem from a Bayesian point of view. In Section 4, we derive an iterative reweighted ℓ_1 lasso algorithm to solve the nonconvex optimisation problem based on concave-convex procedure. In Section 5, we review ADMM, apply it to our optimisation problem, and derive a distributed algorithm. In Section 6, we apply our method to the reconstruction of networks of interconnected Kuramoto oscillators. Finally, in Section 7, we conclude and discuss several future problems.

2. PROBLEM FORMULATION

2.1 Nonlinear Dynamical Systems

We consider dynamical systems described by multi-input multi-output (MIMO) nonlinear discrete-time equations with additive noise:

$$\mathbf{x}(t_{k+1}) = \mathbf{F}(\mathbf{x}(t_k), \mathbf{u}(t_k)) + \boldsymbol{\xi}(t_k), \tag{1}$$

where $\mathbf{x} = [x_1, \ldots, x_{n_\mathbf{x}}]^{\mathrm{T}} \in \mathbb{R}^{n_\mathbf{x}}$ denotes the state vector; $\mathbf{u} = [u_1, \ldots, u_{n_\mathbf{u}}]^{\mathrm{T}} \in \mathbb{R}^{n_\mathbf{u}}$ denotes the input vector; $\mathbf{F}(\cdot) \triangleq [\mathbf{F}_1(\cdot), \ldots, \mathbf{F}_{n_\mathbf{x}}(\cdot)]^T : \mathbb{R}^{n_\mathbf{x}+n_\mathbf{u}} \to \mathbb{R}^{n_\mathbf{x}}$, and $\boldsymbol{\xi} = [\xi_1, \xi_2, \ldots, \xi_{n_\mathbf{x}}]^{\mathrm{T}} \in \mathbb{R}^{n_\mathbf{x}}$ is assumed to be a zero-mean Gaussian white noise vector with constant positive covariance matrix $\boldsymbol{\Xi}$, i.e., $\boldsymbol{\xi}(t_k) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Xi})$. Since this description covers most of the discrete-time nonlinear dynamical systems with infinite number of possible functional forms for $\mathbf{F}(\cdot)$, we confine the scope and assume system (1) satisfies the following assumptions:

Assumption 1. The system (1) is fully measurable, i.e., all the state variables x_i can be measured and there are no hidden variables.

Assumption 2. The function terms $\mathbf{F}(\mathbf{x}(t_k), \mathbf{u}(t_k))$ in (1) are smooth and can be represented as a linear combinations of several dictionary functions, (see Sec. 5.4 in Ljung (1999)).

2.2 Construction of Dictionary Functions

Depending on the field for which the dynamical model needs to be built, only a few typical nonlinearities specific to this field need to be considered. For example, the class of models that arise from genetic regulatory networks (GRN) typically involves nonlinearities that capture fundamental biochemical kinetic laws, e.g., first-order degradation functions, mass-action kinetics, Hill and Michaelis-Menten functions, which are confined to either polynomial or rational functions. In what follows we gather the set of *all* candidate/possible dictionary functions that we want to consider for reconstruction. Consider state variable x_i, $i = 1, \ldots, n_\mathbf{x}$. Under Assumption 2, the function terms for state i can be written as:

$$\begin{aligned}\mathbf{F}_i(\mathbf{x}(t_k), \mathbf{u}(t_k)) &= \sum_{s=1}^{N_i} w_{is} \mathbf{f}_{is}(\mathbf{x}(t_k), \mathbf{u}(t_k)), \\ &= \mathbf{w}_i^{\mathrm{T}} \mathbf{f}_i(\mathbf{x}(t_k), \mathbf{u}(t_k)),\end{aligned} \tag{2}$$

where $\mathbf{w} \in \mathbb{R}^{N_i}$ and $\mathbf{f}_i : \mathbb{R}^{n_\mathbf{x}+n_\mathbf{u}} \to \mathbb{R}^{N_i}$ are dictionary functions that are assumed to govern the dynamics. $\mathbf{f}_i(\mathbf{x}(t_k), \mathbf{u}(t_k))$ can be monomial, polynomial, constant or any other functional form such as rational, exponential, trigonometric etc.

Taking the transpose of both sides of (2), we obtain

$$x_i(t_{k+1}) = \mathbf{f}_i^T(\mathbf{x}(t_k), \mathbf{u}(t_k))\mathbf{w}_i + \xi_i(t_k),\ i = 1, \ldots, n_\mathbf{x}, \tag{3}$$

where $\xi_i(t_k)$ is assumed to be i.i.d. Gaussian distributed: $\xi_i(t_k) \sim \mathcal{N}(0, \sigma_i^2)$, with $\mathbb{E}(\xi_i(t_p)) = 0$, $\mathbb{E}(\xi_i(t_p)\xi_i(t_q)) = \sigma_i^2\delta_{pq}$, with $\delta_{pq} = \begin{cases} 1, & p = q, \\ 0, & p \neq q \end{cases}$. If $M+1$ data samples including the initial value satisfying (3) can be obtained from the system of interest, the system in (3) can be written as

$$\mathbf{y}_i = \mathbf{A}_i \mathbf{w}_i + \boldsymbol{\xi}_i,\ i = 1, \ldots, n_\mathbf{x}. \tag{4}$$

with

$$\begin{aligned}
&\mathbf{y}_i \triangleq [x_i(t_1), \ldots, x_i(t_M)]^{\mathrm{T}} \in \mathbb{R}^M, \\
&\mathbf{A}_i \triangleq \begin{bmatrix} f_{i1}(\mathbf{x}(t_0), \mathbf{u}(t_0)) & \cdots & f_{iN_i}(\mathbf{x}(t_0), \mathbf{u}(t_0)) \\ \vdots & \ddots & \vdots \\ f_{i1}(\mathbf{x}(t_{M-1}), \mathbf{u}(t_{M-1})) & \cdots & f_{iN_i}(\mathbf{x}(t_{M-1}), \mathbf{u}(t_{M-1})) \end{bmatrix} \\
&\quad \in \mathbb{R}^{M \times N_i}, \\
&\mathbf{w}_i \triangleq \begin{bmatrix} w_{i1}, \ldots, w_{iN_i} \end{bmatrix}^T \in \mathbb{R}^{N_i}, \\
&\boldsymbol{\xi}_i \triangleq [\xi_i(t_0), \ldots, \xi_i(t_{M-1})]^{\mathrm{T}} \in \mathbb{R}^M.
\end{aligned}$$

Since the $n_\mathbf{x}$ linear regression problems in (4) are independent, for simplicity of notation, we omit the subscript i in (4) and write

$$\mathbf{y} = \mathbf{A}\mathbf{w} + \boldsymbol{\xi}. \tag{5}$$

The problem is thus to find $\mathbf{w}$ given the measured noisy data stored in $\mathbf{y}$.

2.3 Discussion on relaxations of solutions

Considering the network reconstruction problem in practice, several problems arise with respect to (5). Firstly, a low number of time-series measurements will render the linear regression in (5) under-determined. Secondly, the number of columns of the dictionary matrix might be very large, due to the potential introduction of non-relevant and/or non-independent dictionary functions in $\mathbf{A}$. As a result, the sparsest solution will be favoured due to the model selection criterions such as Akaike information criterion, (Akaike (1974)) or Bayesian information criterion, (Schwarz (1978)). However, finding the sparsest solution is NP-hard. Classically, a lasso algorithm is typically used as a relaxation (Tibshirani (1996)) to this NP-hard optimisation problem. Lasso usually works well when the dictionary matrix has certain properties such as the *restricted isometry property* (RIP), (Candès and Tao (2005) or the incoherence property, (Donoho and Elad (2003)). These properties basically state that two or more of the columns of dictionary matrix cannot be co-linear or close to be co-linear. Unfortunately, such properties are hardly guaranteed in typical network reconstruction problems. In the following, we shall introduce, from a probabilistic viewpoint, how a Bayesian treatment can alleviate these RIP or incoherence requirements (Tipping (2001); Seeger and Wipf (2010)).

3. BAYESIAN VIEWPOINT

3.1 Sparsity Inducing Prior

Bayesian modelling treats all unknowns as stochastic variables with certain probability distributions, (Bishop (2006)). For $\mathbf{y} = \mathbf{A}\mathbf{w} + \boldsymbol{\xi}$, it is assumed that the stochastic variables in $\boldsymbol{\xi}$ are i.i.d. Gaussian distributed with $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$. In such case, the likelihood of the data given $\mathbf{w}$ is

$$\mathcal{P}(\mathbf{y}|\mathbf{w}) = \mathcal{N}(\mathbf{y}|\mathbf{A}\mathbf{w}, \sigma^2\mathbf{I}) \propto \exp\left[-\frac{1}{2\sigma^2}\|\mathbf{A}\mathbf{w} - \mathbf{y}\|_2^2\right]. \tag{6}$$

Given the likelihood function in (6) and specifying a prior $\mathcal{P}(\mathbf{w}) = \prod_j \mathcal{P}(w_j)$. We further define a prior distribution $\mathcal{P}(\mathbf{w})$ as $\mathcal{P}(\mathbf{w}) \propto \exp\left[-\frac{1}{2}g(\mathbf{w})\right] = \exp\left[-\frac{1}{2}\sum_j g(w_j)\right]$, where $g(w_j)$ is a given function of w_j. To enforce sparsity on $\mathbf{w}$, the function $g(\mathbf{w})$ is usually chosen as a concave, non-decreasing function of $|\mathbf{w}|$. Such penalty functions can be realised by adding sparsity inducing priors on $\mathbf{w}$. Such sparsity inducing priors include Laplace $\propto \exp(-\gamma\sum_j |w_j|)$, where $\gamma > 0$, Student's $t \propto (b + w_j^2/2)^{-(a+\frac{1}{2})}$, where $a, b > 0$, and others, (Palmer et al. (2005)).

Typically, $\mathcal{P}(\mathbf{w}|\mathbf{y})$ is approximated by a Gaussian distribution from which efficient algorithms exist (Bishop (2006)). To this end, we may consider *super-Gaussian* priors, which yield a lower bound for the priors $\mathcal{P}(w_j)$ (Palmer et al. (2005)). More specifically, if we define $\boldsymbol{\gamma} \triangleq [\gamma_1, \ldots, \gamma_N]^T \in \mathbb{R}_+^N$, we can represent the prior in the following relaxed (variational) form:

$$\mathcal{P}(\mathbf{w}) = \prod_{j=1}^{n} \mathcal{P}(w_j),\ \mathcal{P}(w_j) = \max_{\gamma_j > 0} \mathcal{N}(w_j|0, \gamma_j)\varphi(\gamma_j), \tag{7}$$

where $\varphi(\gamma_j)$ is a nonnegative function which is treated as a hyperprior with γ_j being its associated hyperparameters. Throughout, we call $\varphi(\gamma_j)$ the "*potential function*". This Gaussian relaxation is possible if and only if $\log \mathcal{P}(\sqrt{w_j})$ is concave on $(0, \infty)$. Hereafter, we adopt Student's t prior and the corresponding $\varphi(\gamma_j) = 1$.

3.2 Marginal Likelihood Maximisation

For a fixed $\boldsymbol{\gamma} = [\gamma_1, \ldots, \gamma_N]$, we define a relaxed prior which is a joint probability distribution over $\mathbf{w}$ and $\boldsymbol{\gamma}$

$$\mathcal{P}(\mathbf{w}; \boldsymbol{\gamma}) = \prod_j \mathcal{N}(w_j|0, \gamma_j)\varphi(\gamma_j) = \mathcal{P}(\mathbf{w}|\boldsymbol{\gamma})\mathcal{P}(\boldsymbol{\gamma}) \leq \mathcal{P}(\mathbf{w}), \tag{8}$$

where $\mathcal{P}(\mathbf{w}|\boldsymbol{\gamma}) \triangleq \prod_j \mathcal{N}(w_j|0, \gamma_j), \mathcal{P}(\boldsymbol{\gamma}) \triangleq \prod_j \varphi(\gamma_j)$. Since is $\mathcal{P}(\mathbf{y}|\mathbf{w})$ is Gaussian in (6), we can get a relaxed posterior which is also Gaussian.

Now the key question is how to choose the most appropriate $\boldsymbol{\gamma} = \hat{\boldsymbol{\gamma}} = [\hat{\gamma}_1, \ldots, \hat{\gamma}_N]$ to maximise $\prod_j \mathcal{N}(w_j|0, \gamma_j)\varphi(\gamma_j)$ such that $\mathcal{P}(\mathbf{w}|\mathbf{y}, \hat{\boldsymbol{\gamma}})$ can be a "good" relaxation to $\mathcal{P}(\mathbf{w}|\mathbf{y})$. Using the product rule for probabilities, we can write the full posterior $\mathcal{P}(\mathbf{w}, \boldsymbol{\gamma}|\mathbf{y}) \propto \mathcal{N}(\mathbf{m}_\mathbf{w}, \boldsymbol{\Sigma}_\mathbf{w}) \times \frac{\mathcal{P}(\mathbf{y}|\boldsymbol{\gamma})\mathcal{P}(\boldsymbol{\gamma})}{\mathcal{P}(\mathbf{y})}$. Since $\mathcal{P}(\mathbf{y})$ is independent of $\boldsymbol{\gamma}$, the quantity

$$\mathcal{P}(\mathbf{y}|\boldsymbol{\gamma})\mathcal{P}(\boldsymbol{\gamma}) = \int \mathcal{P}(\mathbf{y}|\mathbf{w})\mathcal{P}(\mathbf{w}|\boldsymbol{\gamma})\mathcal{P}(\boldsymbol{\gamma})d\mathbf{w}$$

is the prime target for variational methods (Wainwright and Jordan (2008)). This quantity is known as evidence or marginal likelihood. A good way of selecting $\hat{\boldsymbol{\gamma}}$ is to choose it as the minimiser of the sum of the misaligned probability mass, e.g.,

$$\begin{aligned}\hat{\boldsymbol{\gamma}} &= \operatorname*{argmin}_{\boldsymbol{\gamma} \geq \mathbf{0}} \int \mathcal{P}(\mathbf{y}|\mathbf{w})\left|\mathcal{P}(\mathbf{w}) - \mathcal{P}(\mathbf{w}; \boldsymbol{\gamma})\right| d\mathbf{w} \\ &= \arg\max_{\boldsymbol{\gamma} \geq \mathbf{0}} \int \mathcal{P}(\mathbf{y}|\mathbf{w}) \prod_{j=1}^{n} \mathcal{N}(w_j|0, \gamma_j)\varphi(\gamma_j)d\mathbf{w}.\end{aligned}$$

The second equality is a consequence of $\mathcal{P}(\mathbf{w}; \boldsymbol{\gamma}) \leq \mathcal{P}(\mathbf{w})$ (see (8)). The procedure in (9) is referred to as evidence/marginal likelihood maximisation or type-II maximum likelihood, (Tipping (2001); Seeger and Wipf (2010); Wipf et al. (2011); Seeger and Nickisch (2011)). It means that the marginal likelihood can be maximised by selecting the most probable hyperparameters able to explain the observed data. Once $\hat{\boldsymbol{\gamma}}$ is computed, an estimate of the unknown weights can be obtained by setting $\hat{\mathbf{w}}$ to the posterior mean

$$\hat{\mathbf{w}} = \mathbb{E}(\mathbf{w}|\mathbf{y}; \hat{\boldsymbol{\gamma}}) = \hat{\boldsymbol{\Gamma}}\mathbf{A}^T(\sigma^2\mathbf{I} + \mathbf{A}\hat{\boldsymbol{\Gamma}}\mathbf{A}^{\mathrm{T}})^{-1}\mathbf{y}. \tag{9}$$

with $\hat{\boldsymbol{\Gamma}} \triangleq \mathrm{diag}[\hat{\boldsymbol{\gamma}}]$. If an algorithm can be proposed to compute $\hat{\boldsymbol{\gamma}}$ in (9), we can obtain an estimation of the posterior mean $\hat{\mathbf{w}}$ in (9).

In (Tipping (2001); Wipf et al. (2011)), it is shown that $\hat{\boldsymbol{\Gamma}}$ can be obtained by minimising the following nonconvex cost function

$$\mathcal{L}_{\boldsymbol{\gamma}}(\boldsymbol{\gamma}) = \log\left|\sigma^2\mathbf{I} + \mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^{\mathrm{T}}\right| + \mathbf{y}^{\mathrm{T}}(\sigma^2\mathbf{I} + \mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^{\mathrm{T}})^{-1}\mathbf{y}. \tag{10}$$

However, such formulation does not allow to incorporate convex constraints on $\mathbf{w}$, which are typically very useful in network reconstruction. In (Wipf et al. (2011); Pan et al. (2013)), it is shown that the dual cost function in $\mathbf{w}$-space has the following form

$$\min_{\boldsymbol{\gamma} \geq \mathbf{0}, \mathbf{w}} \|\mathbf{A}\mathbf{w} - \mathbf{y}\|_2^2 + \sigma^2\mathbf{w}^T\boldsymbol{\Gamma}^{-1}\mathbf{w} + \log|\sigma^2\mathbf{I} + \mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^{\mathrm{T}}|, \tag{11}$$

To ease notation and avoid interrupting the flow, we will derive the algorithm without considering convex constraints in the sequel.

4. CONCAVE-CONVEX PROCEDURE

The cost function in (19) is convex in $\mathbf{w}$ but nonconvex in $\boldsymbol{\Gamma}$. In the next section, we show how this nonconvex optimisation problem can be formulated as a concave-convex procedure (CCCP) and then, show that finding the solution to CCCP is equivalent to solving iterative reweighted lasso problem. CCCP is another interpretation of the derivation of iterative reweighted lasso algorithm (Wipf and Nagarajan (2010)).

Let

$$u(\mathbf{w}, \boldsymbol{\gamma}) \triangleq \|\mathbf{A}\mathbf{w} - \mathbf{y}\|_2^2 + \sigma^2 \sum_j \frac{w_j^2}{\gamma_j},$$

$$v(\boldsymbol{\gamma}) \triangleq -\log|\sigma^2\mathbf{I} + \mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^{\mathrm{T}}|$$

Note that $u(\mathbf{w}, \boldsymbol{\gamma})$ is jointly convex in $\mathbf{w}$ and $\boldsymbol{\gamma}$, $v(\boldsymbol{\gamma})$ is convex in $\boldsymbol{\gamma}$. Then the minimisation of the cost function (19) can be formulated as a concave-convex procedure

$$\min_{\boldsymbol{\gamma} \geq \mathbf{0}, \mathbf{w}} u(\mathbf{w}, \boldsymbol{\gamma}) - v(\boldsymbol{\gamma}) \tag{12}$$

Since $v(\boldsymbol{\gamma})$ is differentiable over $\boldsymbol{\gamma}$, the problem in (12) can be transformed into the following iterative convex optimisation problem

$$\left[\mathbf{w}^{k+1}, \boldsymbol{\gamma}^{k+1}\right] = \underset{\boldsymbol{\gamma}\geq\mathbf{0},\mathbf{w}}{\operatorname{argmin}}\, u(\mathbf{w},\boldsymbol{\gamma}) - \nabla_{\boldsymbol{\gamma}} v(\boldsymbol{\gamma}^k)^T \boldsymbol{\gamma}. \quad (13)$$

Using basic principles in convex analysis, we then obtain the following analytic form for the negative gradient of $v(\boldsymbol{\gamma})$ at $\boldsymbol{\gamma}$is:

$$\begin{aligned}\boldsymbol{\alpha}^k &= -\nabla_{\boldsymbol{\gamma}} v(\boldsymbol{\gamma}^k)^T \\ &= -\nabla_{\boldsymbol{\gamma}} \left(-\log|\sigma^2\mathbf{I} + \mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^{\mathrm{T}}|\right)|_{\boldsymbol{\gamma}=\boldsymbol{\gamma}^k} \\ &= \operatorname{diag}\left[\mathbf{A}^{\mathrm{T}}\left(\sigma^2\mathbf{I} + \mathbf{A}\boldsymbol{\Gamma}^{\mathbf{k}}\mathbf{A}^{\mathrm{T}}\right)^{-1}\mathbf{A}\right],\end{aligned}$$

Then the iterative procedure (13) can be formulated as

$$\left[\mathbf{w}^{k+1}, \boldsymbol{\gamma}^{k+1}\right] = \underset{\boldsymbol{\gamma}\geq\mathbf{0},\mathbf{w}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{w}-\mathbf{y}\|_2^2 + \sigma^2 \sum_j \left(\frac{w_j^2}{\gamma_j} + \alpha_j^k \gamma_j\right). \quad (14)$$

The objective function in (14) is jointly convex in $\mathbf{w}$ and $\boldsymbol{\gamma}$ and can be globally minimised by solving over $\boldsymbol{\gamma}$ and then $\mathbf{w}$. If $\mathbf{w}$ is fixed, it gives

$$\boldsymbol{\gamma}^{k+1} = \underset{\boldsymbol{\gamma}\geq\mathbf{0}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{w}-\mathbf{y}\|_2^2 + \sigma^2 \sum_j \left(\frac{w_j^2}{\gamma_j} + \alpha_j^k \gamma_j\right). \quad (15)$$

We notice that in (15), $\boldsymbol{\gamma}^{k+1}$ has closed form solution $\gamma_j^{k+1} = |w_j|/\sqrt{\alpha_j^k}$, If $\gamma_j^{k+1} = |w_j|/\sqrt{\alpha_j^k}$ is substituted into (15), we get

$$\begin{aligned}\mathbf{w}^{k+1} &= \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{w}-\mathbf{y}\|_2^2 + \sigma^2 \sum_j \left(\frac{w_j^2}{\gamma_j^{k+1}} + \alpha_j^k \gamma_j^{k+1}\right) \\ &= \underset{\mathbf{w}}{\operatorname{argmin}} \{\|\mathbf{A}\mathbf{w}-\mathbf{y}\|_2^2 + 2\sigma^2 \sum_{j=1}^N \sqrt{\alpha_j^k}|w_j|\}.\end{aligned}$$

We can then set $\gamma_j^{k+1} = |w_j^{k+1}|/\sqrt{\alpha_j^k}, \forall j$. Then we update $\boldsymbol{\alpha}^{k+1}$ by (14). However, some of the estimated weights will be several orders of magnitude lower than the average "energy", e.g., $w_j^2 \ll \|\mathbf{w}\|_2^2$. Thus a threshold needs to be defined *a priori* to prune "small" weights at each iteration.

We can now explain how the update of the parameters can be performed based on the above. Set the iteration count k to zero and $\theta_j^0 = 1,\ \forall j$. At this stage, the optimisation is a typical lasso. Then at the k^{th} iteration, we initialise $\theta_j^{(k)} = \sqrt{\alpha_j^k}, \forall j$. The above described procedure is summarised in Algorithm 1.

In the next section, we will reformulate the centralised optimisation in Algorithm 1 into distributed optimisation by ADMM.

Remark 1. There is one interesting finding from (15). We can actually minimise over $\mathbf{w}$ first then minimise over $\boldsymbol{\gamma}$. Rather than only initialise $\boldsymbol{\alpha}^0$, we also initialise $\boldsymbol{\gamma}^0$. The iterative procedure is then a reweighted ℓ_2 algorithm (Wipf and Nagarajan (2010))

$$\mathbf{w}^{k+1} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{w}-\mathbf{y}\|_2^2 + \sigma^2 \sum_j \frac{w_j^2}{\gamma_j^k}. \quad (17)$$

Algorithm 1 Reweighted lasso on $\mathbf{w}$

1: Initialise $\theta_j^0 = 1,\ \forall j$
2: **for** $k = 0, \ldots, k_{\max}$ **do**
3: Solve the weighted lasso problem

$$\mathbf{w}^{k+1} = \underset{\mathbf{w}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{A}\mathbf{w}-\mathbf{y}\|_2^2 + \sigma^2\|\boldsymbol{\Theta}^k\mathbf{w}\|_1; \quad (16)$$

4: Set $\boldsymbol{\Theta}^{(k)} \triangleq \operatorname{diag}\left[\boldsymbol{\theta}^{(k)}\right]^{-1}$, $\mathbf{W}^{(k)} \triangleq \operatorname{diag}\left[|\mathbf{w}^{(k)}|\right]$ and $\theta_j^{k+1} = \left[(\mathbf{A}_j)^{\mathrm{T}}\left(\sigma^2\mathbf{I} + \mathbf{A}(\boldsymbol{\Theta}^k)^{-1}\mathbf{W}^{k+1}(\mathbf{A})^T\right)^{-1}\mathbf{A}_j\right]^{\frac{1}{2}}$
5: **if** A stopping criterion is satisfied **then**
6: Break;
7: **end if**
8: **end for**

Then with the estimate $\mathbf{w}^{k+1}$

$$\boldsymbol{\gamma}^{k+1} = \underset{\boldsymbol{\gamma}\geq\mathbf{0}}{\operatorname{argmin}} \sum_j \left(\frac{(w_j^{k+1})^2}{\gamma_j} + \alpha_j^k \gamma_j\right). \quad (18)$$

$\boldsymbol{\gamma}^{k+1}$ has a closed form solution $\boldsymbol{\gamma}^{k+1} = |w_j^{k+1}|/\sqrt{\alpha_j^k}$. Then we can update $\boldsymbol{\alpha}^{k+1}$ as in (14) and iterate until a stopping criterion is satisfied.

The $\mathbf{w}$ update in (17) is actually a reweighted-ℓ_2 algorithm. $\boldsymbol{\gamma}$ are the weighting vectors like $\boldsymbol{\theta}$ in reweighted lasso in Algorithm 1. We also tested this reweighted ℓ_2 algorithm for the examples in the sequel. However the heuristic convergence rate and reconstruction accuracy is not as good as reweighted lasso algorithm. We are still studying on it.

5. ALTERNATING DIRECTION METHOD OF MULTIPLIERS (ADMM)

In this section we give an overview of ADMM. We follow closely the development of (Boyd et al. (2011)

5.1 ADMM

ADMM is a numerical algorithm for solving optimisation problems such as

$$\begin{aligned}\min_{\mathbf{w}} \quad & f(\mathbf{w}) + g(\mathbf{z}), \\ \text{subject to} \quad & P\mathbf{w} + Q\mathbf{z} = \mathbf{c},\end{aligned} \quad (19)$$

with variable $\mathbf{w} \in \mathbb{R}^n$ and $\mathbf{z} \in \mathbb{R}^m$, where $P \in \mathbb{R}^{p\times n}$, $Q \in \mathbb{R}^{p\times m}$, and $\mathbf{c} \in \mathbb{R}^p$. We will assume that f and g are convex.

As the method of multipliers, we form the augmented Lagrangian

$$\begin{aligned}L_\rho =& f(\mathbf{w}) + g(\mathbf{z}) + \mathbf{v}^T(P\mathbf{w} + Q\mathbf{z} - c) + \\ & \rho/2\|P\mathbf{w} + Q\mathbf{z} - c\|_2^2.\end{aligned} \quad (20)$$

Defining the residual $r = P\mathbf{w} + Q\mathbf{z} - c$ and $\mathbf{u} = (1/\rho)\mathbf{v}$ as the scaled dual variable, we can express ADMM as

$$\begin{aligned}\mathbf{w}^{k+1} &:= \underset{\mathbf{w}}{\operatorname{argmin}} \left(f(\mathbf{w}) + \frac{\rho}{2}\|P\mathbf{w} + Q\mathbf{z}^k - \mathbf{c} + \mathbf{u}^k\|_2^2\right) \\ \mathbf{z}^{k+1} &:= \underset{\mathbf{z}}{\operatorname{argmin}} \left(g(\mathbf{z}) + \frac{\rho}{2}\|P\mathbf{w}^{k+1} + Q\mathbf{z} - \mathbf{c} + \mathbf{u}^k\|_2^2\right) \\ \mathbf{u}^{k+1} &:= \mathbf{u}^k + P\mathbf{w}^{k+1} + Q\mathbf{z}^{k+1} - \mathbf{c}.\end{aligned}$$

Stopping criterion The primal and dual residuals at iteration k are given by

$$\mathbf{e}^k_{primal} = \mathbf{w}^k - \mathbf{z}^k, \mathbf{e}^k_{dual} = -\rho(\mathbf{z}^k - \mathbf{z}^{k-1}).$$

We terminate the algorithm when the primal and dual residuals satisfy a stopping criterion:

$$\|\mathbf{e}^k_{primal}\|_2 \leq \epsilon_{primal}, \|\mathbf{e}^k_{dual}\|_2 \leq \epsilon_{dual}.$$

Here, the tolerances $\epsilon_{primal} > 0$ and $\epsilon_{dual} > 0$ and can be set via an absolute plus relative criterion

$$\epsilon_{primal} = \sqrt{n}\epsilon_{abs} + \epsilon_{rel}\max(\|\mathbf{w}^k\|_2, \|\mathbf{z}^k\|_2),$$
$$\epsilon_{dual} = \sqrt{n}\epsilon_{abs} + \epsilon_{rel}\rho\|\mathbf{u}^k\|,$$

where ϵ_{abs} and ϵ_{rel} are absolute and relative tolerances. More details can be found in (Boyd et al. (2011)).

5.2 Splitting across candidate functions

In our setting, the number of candidate functions will be very large. Therefore we partition across the candidate functions. Each subsystem can deal with its split of candidate functions independently then update the shared variables. The following are direct consequences of the so-called sharing problem in (Boyd et al. (2011)).

We partition the parameter vector $\mathbf{w}$ as $\mathbf{w} = (\mathbf{w}_1, \ldots, \mathbf{w}_n)$, with $\mathbf{w}_i \in \mathbb{R}^{N_i}$, where $\sum_{i=1}^n N_i = N$. Partition the dictionary matrix $\mathbf{A}$ as $\mathbf{A} = [\mathbf{A}_1, \ldots, \mathbf{A}_n]$, with $\mathbf{A}_i \in \mathbb{R}^{M\times N_i}$. Thus $\mathbf{A}\mathbf{w} = \sum_{i=1}^n \mathbf{A}_i\mathbf{w}_i$, i.e., $\mathbf{A}_i\mathbf{w}_i$ can be thought of as a 'partial' prediction of $\mathbf{y}$ using only the candidate functions referenced in $\mathbf{w}_i$.

Then the reweighted lasso problem (16)

$$\min_{\mathbf{w}} \frac{1}{2}\|\mathbf{A}\mathbf{w} - \mathbf{y}\|_2^2 + \lambda\|\mathbf{\Theta}\mathbf{w}\|_1$$

becomes

$$\min_{\mathbf{w}} \frac{1}{2}\|\sum_{i=1}^N \mathbf{A}_i\mathbf{w}_i - \mathbf{y}\|_2^2 + \lambda\sum_{i=1}^N \|\mathbf{\Theta}_i\mathbf{w}_i\|_1.$$

Following the approach used for the sharing problem, we express the problem as

$$\begin{aligned} \min \quad & \frac{1}{2}\|\sum_{i=1}^N \mathbf{z}_i - \mathbf{y}\|_2^2 + \lambda\sum_{i=1}^N \|\mathbf{\Theta}_i\mathbf{w}_i\|_1, \\ \text{subject to} \quad & \mathbf{A}_i\mathbf{w}_i - \mathbf{z}_i = 0,\ i = 1, \ldots, N, \end{aligned}$$

with new variables $\mathbf{z} \in \mathbb{R}^m$. The derivation and simplification of ADMM also follows that for the sharing problem. The scaled form of ADMM is

$$\begin{aligned} \mathbf{w}_i^{k+1} &:= \operatorname*{argmin}_{\mathbf{w}_i} \frac{\rho}{2}\|\mathbf{A}_i\mathbf{w}_i - \mathbf{z}^k + \mathbf{u}^k\|_2^2 + \lambda\|\mathbf{\Theta}_i\mathbf{w}_i\|_1 \\ \mathbf{z}^{k+1} &:= \operatorname*{argmin}_{\mathbf{z}} \frac{1}{2}\|\sum_{i=1}^N \mathbf{z}_i - \mathbf{y}\|_2^2 + \sum_{i=1}^N \frac{\rho}{2}\|\mathbf{A}\mathbf{w}^{k+1} - \mathbf{z}^k + \mathbf{u}^k\|_2^2 \\ \mathbf{u}^{k+1} &:= \mathbf{u}^k + \mathbf{A}\mathbf{w}^{k+1} - \mathbf{z}^{k+1}. \end{aligned} \tag{21}$$

As in the discussion for the sharing problem, we carry out the $\mathbf{z}$-update by first solving for the average

$$\begin{aligned} \overline{\mathbf{z}}^{k+1} &:= \operatorname*{argmin}_{\overline{\mathbf{z}}} \|N\overline{\mathbf{z}} - \mathbf{y}\|_2^2 + \sum_{i=1}^N \frac{N\rho}{2}\|\overline{\mathbf{z}} - \overline{\mathbf{A}\mathbf{w}}^{k+1} - \overline{\mathbf{u}}^k\|_2^2 \\ \mathbf{z}^{k+1} &:= \overline{\mathbf{z}}^{k+1} + \mathbf{A}_i\mathbf{w}_i^{k+1} + \mathbf{u}^k - \overline{\mathbf{A}\mathbf{w}}^{k+1} - \overline{\mathbf{u}}^k, \end{aligned} \tag{22}$$

where $\overline{\mathbf{A}\mathbf{w}}^{k+1} = (1/N)\sum_{i=1}^N \mathbf{A}_i\mathbf{w}_i^{k+1}$. Substituting the last expression into the update for $\mathbf{u}_i$, we find that

$$\mathbf{u}_i^{k+1} = \overline{\mathbf{A}\mathbf{w}}^{k+1} + \overline{\mathbf{u}}^k - \overline{\mathbf{z}}^{k+1},$$

which shows that, as in the sharing problem, all the dual variables are equal. Using a single dual variable $\mathbf{u}^k \in \mathbb{R}^m$, eliminating $\mathbf{z}_i$, and define

$$\mathbf{b} = \mathbf{A}_i\mathbf{w}_i^k + \overline{\mathbf{z}}^k - \overline{\mathbf{A}\mathbf{w}}^k - \mathbf{u}^k$$

we arrive at the Algorithm 2

Algorithm 2 ADMM for splitting across candidate functions

1: **for** $k = 0, \ldots, k_{\max}$ **do**
2: Solve
$$\mathbf{w}_i^{k+1} = \operatorname*{argmin}_{\mathbf{w}_i} \left(\frac{\rho}{2}\|\mathbf{A}_i\mathbf{w}_i - \mathbf{b}\|_2^2 + \lambda\|\mathbf{\Theta}_i\mathbf{w}_i\|_1\right);$$
3: Compute $\overline{\mathbf{z}}^{k+1} = \frac{1}{N+\rho}\left(\mathbf{y} + \rho\overline{\mathbf{A}\mathbf{w}}^{k+1} + \rho\mathbf{u}^k\right)$;
4: Update $\mathbf{u}^{k+1} = \mathbf{u}^k + \overline{\mathbf{A}\mathbf{w}}^{k+1} - \overline{\mathbf{z}}^{k+1}$;
5: **if** A stopping criterion is satisfied **then**
6: Break;
7: **end if**
8: **end for**

Each $\mathbf{w}_i$-update is a lasso problem with n_i variables, which can be solved using any lasso method.

In the $\mathbf{w}_i$-update, we have $\mathbf{w}_i^{k+1} = \mathbf{0}$ (meaning that none of the features in the i-th block are used) if and only if

$$\|\mathbf{A}_i^T\mathbf{b}\|_2 \leq \frac{\lambda}{\rho} \tag{23}$$

The first step involves solving N parallel weighted lasso (weighted ℓ_1-regularised least squares) problems in n_i variables each. Between the first and second steps, we collect and sum the partial predictors $\mathbf{A}_i\mathbf{w}_i^{k+1}$ to form $\overline{\mathbf{A}\mathbf{w}}^{k+1}$. The second step is a single minimisation in m variables, a quadratically regularised loss minimisation problem

5.3 ADMM for Weighted Lasso

The weighted lasso problem can be solved using ADMM as well.

$$\begin{aligned} \min \quad & \frac{1}{2}\|\mathbf{A}_i\mathbf{w}_i - \mathbf{b}\|_2^2 + \frac{\lambda}{\rho}\|\mathbf{\Theta}_i\mathbf{w}_i\|_1, \\ \text{subject to} \quad & \mathbf{\Theta}_i\mathbf{w}_i - \hat{\mathbf{z}}_i = 0,\ i = 1, \ldots, N. \end{aligned}$$

Define $\hat{\lambda} = \lambda/\rho$, it yields the ADMM algorithm

Algorithm 3 ADMM for weighted lasso

1: **for** $k = 0, \ldots, k_{\max}$ **do**
2: Update $\mathbf{w}_i^{k+1} = (\mathbf{A}_i^T\mathbf{A}_i + \hat{\rho}\mathbf{\Theta}_i^T\mathbf{\Theta}_i)^{-1}(\mathbf{A}_i^T\mathbf{b} + \hat{\rho}\mathbf{\Theta}_i^T(\hat{\mathbf{z}}_i - \hat{\mathbf{u}}_i))$;
3: Update $\hat{\mathbf{z}}_i^{k+1} = S_{\hat{\lambda}/\hat{\rho}}(\mathbf{\Theta}_i\mathbf{w}_i^{k+1} + \hat{\mathbf{u}}_i^k)$;
4: Update $\hat{\mathbf{u}}_i^{k+1} = \hat{\mathbf{u}}_i^k + \mathbf{\Theta}_i\mathbf{w}_i^{k+1} - \hat{\mathbf{z}}_i^{k+1}$;
5: **if** A stopping criterion is satisfied or when k reaches a predefined iteration number $k_{\max}$ **then**
6: Break;
7: **end if**
8: **end for**

where $\hat{\rho}$ is the penalty parameter and the soft thresholding operator $S_{\hat{\lambda}/\hat{\rho}}$ is defined as

$$S_{\hat{\lambda}/\hat{\rho}}(x) = \max(0, x - \hat{\lambda}/\hat{\rho}) - \max(0, -x - \hat{\lambda}/\hat{\rho}).$$

5.4 Algorithm For Nonlinear Network Reconstruction

Now, we summarise the procedure for nonlinear network reconstruction in Algorithm 4.

Algorithm 4 ADMM on $\mathbf{w}$

1: Initialisation
   (1) Collect the time-series data, specify the candidate functions and construct the dictionary matrix;
   (2) Partition the dictionary matrix $\mathbf{A}$ as $\mathbf{A} = [\mathbf{A}_1, \ldots, \mathbf{A}_P]$, with $\mathbf{A}_i \in \mathbb{R}^{M \times P_i}$;
   (3) Initialise the weight $\mathbf{\Theta}^0$ as $\mathbf{\Theta}^0 = [\mathbf{\Theta}_1^0, \ldots, \mathbf{\Theta}_P^0]$, $\mathbf{\Theta}_i^0 = [\theta_{i1}^0, \ldots, \theta_{iP_i}^0]$, with $\theta_{ij}^0 = 0$.
2: **for** $k = 0, \ldots, k_{\max}$ **do**
3: Apply Algorithm 2 and 3 to to get an estimate on $\mathbf{w}^{k+1}$ to (16);
4: Set $\mathbf{\Theta}^{(k)} \triangleq \operatorname{diag}\left[\boldsymbol{\theta}^{(k)}\right]^{-1}$, $\mathbf{W}^{(k)} \triangleq \operatorname{diag}\left[|\mathbf{w}^{(k)}|\right]$;
5: Update $\theta_j^{(k+1)}$ for the next iteration

$$\theta_j^{(k+1)} = \left[\mathbf{A}_j^{\mathrm{T}} \left(\sigma^2 \mathbf{I} + \mathbf{A}\mathbf{\Theta}^{(k)}\mathbf{W}^{(k+1)}\mathbf{A}^T\right)^{-1} \mathbf{A}_j)\right]^{\frac{1}{2}};$$

6: **if** A stopping criterion is satisfied **then**
7: Break;
8: **end if**
9: **end for**

6. NUMERICAL ILLUSTRATION

6.1 An Example of Kuramoto Oscillator

A classical example in physics, engineering and biology is the Kuramoto oscillator network (Strogatz (2000)). We consider a network where the Kuramoto oscillator are non-identical (each has its own natural oscillation frequency ω_i) and the coupling strengths between nodes are not the same. The corresponding discrete-time dynamics can be described by

$$\phi_i(t_{k+1}) = \phi_i(t_k) + (t_{k+1} - t_k) \left[\omega_i + \sum_{j=1}^{n} w_{ij} g_{ij}(\phi_j(t_k) - \phi_i(t_k)) + \xi_i(t_k)\right], \tag{24}$$

where $i = 1, \ldots, n$, $\phi_i \in [0, 2\pi)$ is the phase of oscillator i, ω_i is its natural frequency, and the coupling function g_{ij} is usually taken as sine for all i, j. w_{ij} represent the coupling strength between oscillators i and j thus $[w_{ij}]_{n\times n}$ defines the topology of the oscillator network. Here, assuming we don't know the exact form of g_{ij}, we reconstruct from time-series data of the individual phases ϕ_i a dynamical network consisting of n Kuramoto oscillators, i.e., we identify the coupling functions $g_{ij}(\cdot)$ as well as the model parameters, i.e., ω_i and w_{ij}, $i, j = 1, \ldots, n$.

To define the dictionary matrix $\mathbf{A}$, we assume that all the dictionary functions are functions of a pair of state variables only and consider 2 candidate coupling functions g_{ij}: $\sin(x_j - x_i)$, $\cos(x_j - x_i)$. Based on this, we define the dictionary matrix as

$$\mathbf{A}_{ij}(x_j(t_k), x_i(t_k)) \triangleq [\sin(x_j(t_k) - x_i(t_k)), \cos(x_j(t_k) - x_i(t_k))] \in \mathbb{R}^2.$$

To also take into account the natural frequencies, we add to the last column of $\mathbf{A}_i$ a unit vector. This leads to the following dictionary matrix $\mathbf{A}_i$:

$$\mathbf{A}_i \triangleq \begin{bmatrix} \mathbf{A}_{i1}(x_1(t_0), x_i(t_0)) & \ldots & \mathbf{A}_{in}(x_n(t_0), x_i(t_0)) & 1 \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{A}_{i1}(x_1(t_{M-1}), x_i(t_{M-1})) & \ldots & \mathbf{A}_{in}(x_n(t_{M-1}), x_i(t_{M-1})) & 1 \end{bmatrix} \in \mathbb{R}^{M \times (2n+1)}.$$

Then the output can be defined as

$$\mathbf{y}_i \triangleq \left[\frac{\phi_i(t_1) - \phi_i(t_0)}{t_1 - t_0}, \ldots, \frac{\phi_i(t_M) - \phi_i(t_{M-1})}{t_M - t_{M-1}}\right]^{\mathrm{T}} \in \mathbb{R}^M.$$

To generate the time-series data, we simulated a Kuramoto oscillator network for which 10% of the non-diagonal entries of the weight matrix $[w_{ij}]_{n\times n}$ are nonzero (assuming g_{ii} and w_{ii} are zeros), and the non-zero w_{ij} values are drawn from a standard uniform distribution on the interval $[-10, 10]$. The natural frequencies ω_i are drawn from a normal distribution with mean 0 and variance 10. In order to create simulated data, we simulated the discrete-time model (24) and took 'measurements data points' every $t_{k+1} - t_k = 0.1$ between $t = 0$ and $t = 100$ (in arbitrary units) from random initial conditions which are drawn from a standard uniform distribution on the open interval $(0, 2\pi)$. Thus a total of 1001 measurements for each oscillator phase ϕ_i are collected (including the initial value). Once again, it should be noted that the the number of rows is less than that of columns of the dictionary matrix.

6.2 Algorithmic Performance Comparisons

The reweighted lasso algorithm can be implemented in a centralised way by using CVX (Grant et al. (2008)) or YALMIP (Lofberg (2004)), Matlab packages for specifying and solving convex optimisation problems. CVX or YALMIP calls generic SDP solvers (SDPT3 (Toh et al. (1999) or SeDuMi (Sturm (1999)) to solve the problem. While these solvers are reliable for wide classes of optimisation problems, they are not customised for particular problem families, such as ours.

We compare the centralised algorithm using CVX, the centralised algorithm using ADMM, and the distributed algorithm using ADMM. We fixed the number of measurements M to be 1001 and varied the network size n between 500 and 100,000. For the distributed algorithm, we split the problem into 1000 subproblems where each one has the same dimension. The algorithm is implemented in MATLAB R2012b. The calculations were performed on HP workstation with two 8 core Intel® Xeon(R) CPU E5-2650 2.00GHz with 64 GB RAM.

Since the reconstruction problem in (4) for each node is independent, we therefore consider the performance of a single node for illustration. We first investigate the performance for different signal-to-noise ratios of the data generated for this example. We define signal-to-noise ratio (SNR) by $\mathrm{SNR(dB)} \triangleq 20\log(\|\mathbf{A}\mathbf{w}_{\mathrm{true}}\|_2/\|\boldsymbol{\xi}\|_2)$. We considered SNR ranging from 5 dB to 25 dB for each generated weight. To characterise the accuracy of a reconstruction, we use the normalised mean square error (NMSE) as a performance index, defined by $\|\hat{\mathbf{w}} - \mathbf{w}\|/\|\mathbf{w}\|$, where $\hat{\mathbf{w}}$

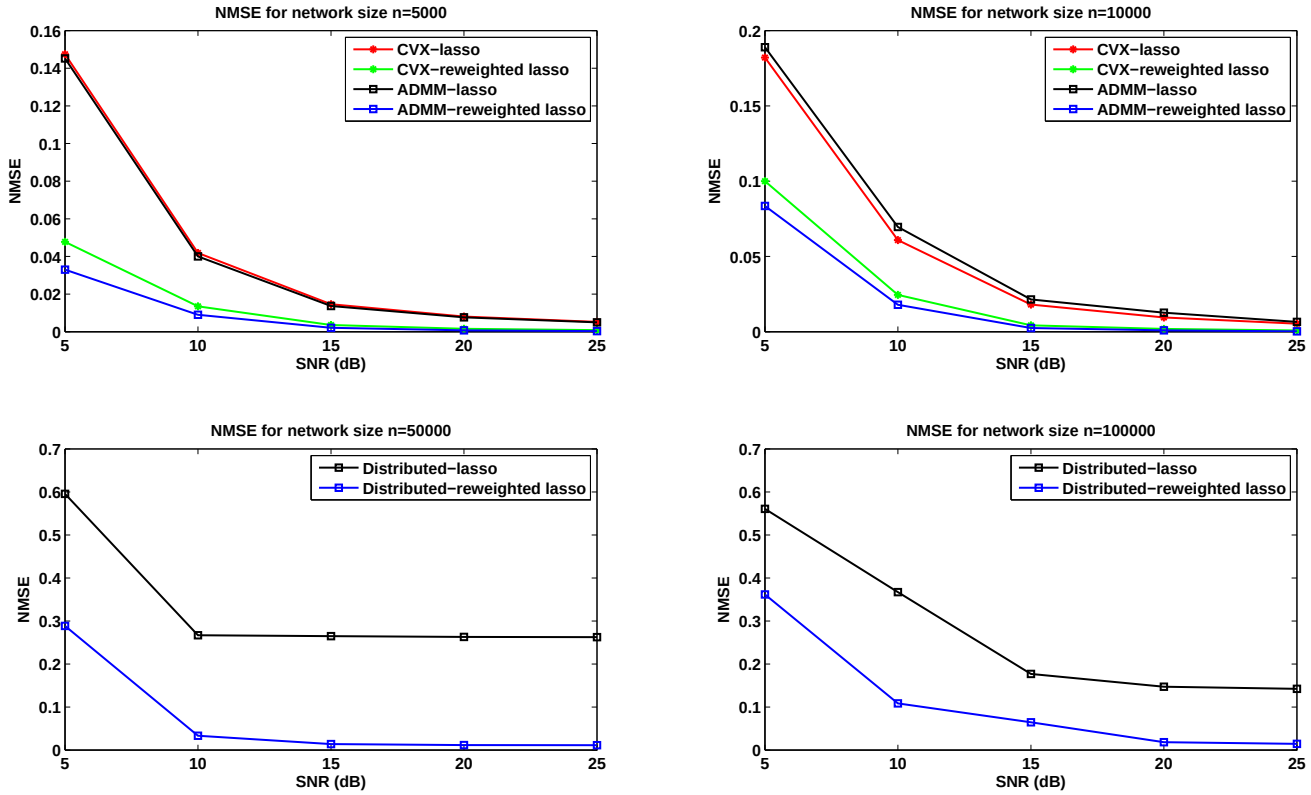

Fig. 1. Normalised Mean Square Error (NMSE) averaged over 50 independent experiments for the signal-to-noise ratios 5dB, 10 dB, 15 dB, 20 dB, and 25 dB, with different network dimensions 5000, 10000, 50000 and 100000.

is the estimate of the true weight $\mathbf{w}$. For each SNR, we generated 50 independent experiments (with different initial conditions and parameters) and calculated the average NMSE for each SNR over these 10 experiments. The results are shown in Fig. 1.

Next, we compare the computation time for different network sizes for each method. These methods are tested under different SNRs as above. In the implementation of the distributed algorithm, we use Matlab command `parfor` to parallelise the $\mathbf{w}_k$-update in Algorithm 2. We use Matlab command `matlabpool('size')` to start a worker pool. The `size` varies from 2 to 10. For each method, we found that the computation time over each SNR varied slightly (at least within the same magnitude). We calculated the average computation time from a total 250 ($=5\times50$) independent experiments for each method. The results are shown in Table 1 (unit in second). For larger problem, with network size 50,000 and 100,000, CVX based reweighted lasso run into memory difficulties. Thus there are no results reported in the table for these network sizes. For the distributed reweighted lasso algorithm, the computation time decreases when the problem is split between an increasing number of processors. However, it should be noted the computation time for distributed reweighted lasso is small partially because Matlab performs some matrix computations in parallel. On the other hand, CVX exploits only one core.

For all the experiments, we set the penalty parameter $\rho = 1$ in the augmented Lagrangian and the scalar regularisation parameter $\lambda = 0.05\|\mathbf{A}^T\mathbf{y}\|_\infty$. We also considered termination tolerances $\epsilon^{abs} = 10^{-4}$ and $\epsilon^{abs} = 10^{-2}$ and set the ADMM iteration number to be 200 and the reweighted iteration number to be 10 throughout all algorithms.

7. CONCLUSION AND DISCUSSION

In this paper, a new distributed reconstruction method for nonlinear dynamical network is proposed. The proposed method only requires time-series data and some prior knowledge about the class of systems for which a dynamical model needs to be built. The network reconstruction problem can be casted as a sparse linear regression problem. Under Bayesian interpretation, this problem is solved using a reweighed ℓ_1 algorithm which can further reduce the Normalised Mean Square Error in comparison with the classic lasso algorithm. Furthermore, our distributed algorithm can deal with networks comprising more than 50,000 nodes, which centralised algorithm typically cannot deal with.

Although the convergence of ADMM algorithm and reweighted lasso algorithm have been studied previously (Boyd et al. (2011); Wipf and Nagarajan (2007)), the convergence of the algorithms developed in this paper still need to be properly characterised. We are currently establishing such convergence results as well as the associated convergence rates and their dependence on parameters such as ρ and λ.

8. ACKNOWLEDGEMENT

Authors would like to thank Dr Ye Yuan and Dr Jorge Gonçalves for helpful discussions and suggestions.

Table 1.

Network Size n / Methods	500	5000	10000	50000	100000
CVX-reweighted lasso	81.9	428.4	1218.3	N/A	N/A
Distributed reweighted lasso with 2 cores	24.7	84.3	156.5	587.1	1341.5
Distributed reweighted lasso with 4 cores	16.4	64.1	91.5	411.8	868.7
Distributed reweighted lasso with 10 cores	15.5	46.9	62.9	345.2	788.3

REFERENCES

H. Akaike. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6): 716–723, 1974.

C.M. Bishop. *Pattern recognition and machine learning*, volume 4. springer New York, 2006.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

E.J. Candès and T. Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12): 4203–4215, 2005.

D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.

D.L. Donoho and M. Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.

Mário AT Figueiredo and José M Bioucas-Dias. Restoration of poissonian images using alternating direction optimization. *Image Processing, IEEE Transactions on*, 19(12):3133–3145, 2010.

M. Grant, S. Boyd, and Y. Ye. Cvx: Matlab software for disciplined convex programming. *Online accessiable: http://stanford. edu/boyd/cvx*, 2008.

R. Haber and H. Unbehauen. Structure identification of nonlinear dynamic systems: survey on input/output approaches. *Automatica*, 26(4):651–677, 1990.

L. Ljung. *System Identification: Theory for the User*. Prentice Hall, 1999.

L. Ljung, H. Hjalmarsson, and H. Ohlsson. Four encounters with system identification. *European Journal of Control*, 17(5):449, 2011.

Johan Lofberg. Yalmip: A toolbox for modeling and optimization in matlab. In *Computer Aided Control Systems Design, 2004 IEEE International Symposium on*, pages 284–289. IEEE, 2004.

Jason Palmer, Kenneth Kreutz-delgado, Bhaskar D Rao, and David P Wipf. Variational em algorithms for non-gaussian latent variable models. In *Advances in neural information processing systems*, pages 1059–1066, 2005.

Wei Pan, Ye Yuan, Jorge Gonçalves, and Guy-Bart Stan. Bayesian approaches to nonlinear network reconstruction. *submitted*, 2013.

G. Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.

Matthias W Seeger and Hannes Nickisch. Large scale bayesian inference and experimental design for sparse linear models. *SIAM Journal on Imaging Sciences*, 4 (1):166–199, 2011.

M.W. Seeger and D.P. Wipf. Variational bayesian inference techniques. *Signal Processing Magazine, IEEE*, 27 (6):81–91, 2010.

J. Sjöberg, Q. Zhang, L. Ljung, A. Benveniste, B. Delyon, P.Y. Glorennec, H. Hjalmarsson, and A. Juditsky. Nonlinear black-box modeling in system identification: a unified overview. *Automatica*, 31(12):1691–1724, 1995.

S.H. Strogatz. From kuramoto to crawford: exploring the onset of synchronisation in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1): 1–20, 2000.

Jos F Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653, 1999.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

M.E. Tipping. Sparse bayesian learning and the relevance vector machine. *The Journal of Machine Learning Research*, 1:211–244, 2001.

Kim-Chuan Toh, Michael J Todd, and Reha H Tütüncü. Sdpt3–a matlab software package for semidefinite programming, version 1.3. *Optimization Methods and Software*, 11(1-4):545–581, 1999.

Bo Wahlberg, Stephen Boyd, Mariette Annergren, and Yang Wang. An admm algorithm for a class of total variation regularized estimation problems. *16th IFAC Symposium on System Identification*, 2012.

M.J. Wainwright and M.I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.

D. Wipf and S. Nagarajan. Iterative reweighted l1 and l2 methods for finding sparse solutions. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):317–329, 2010.

David P Wipf and Srikantan S Nagarajan. A new view of automatic relevance determination. In *Advances in Neural Information Processing Systems*, pages 1625–1632, 2007.

D.P. Wipf, B.D. Rao, and S. Nagarajan. Latent variable bayesian models for promoting sparsity. *Information Theory, IEEE Transactions on*, 57(9):6236–6255, 2011.